

Nonverbal Robot Feedback for Human Teachers

Sandy H. Huang^{*†}, Isabella Huang^{*}, Ravi Pandya^{*}, Anca D. Dragan
University of California, Berkeley

Abstract: Robots can learn preferences from human demonstrations, but their success depends on how informative these demonstrations are. Being informative is unfortunately very challenging, because during teaching, people typically get no transparency into what the robot already knows or has learned so far. In contrast, human students naturally provide a wealth of nonverbal feedback that reveals their level of understanding and engagement. In this work, we study how a robot can similarly provide feedback that is minimally disruptive, yet gives human teachers a better mental model of the robot learner, and thus enables them to teach more effectively. Our idea is that at any point, the robot can indicate what it thinks the correct next action is, shedding light on its current estimate of the human’s preferences. We analyze how useful this feedback is, both in theory and with two user studies—one with a virtual character that tests the feedback itself, and one with a PR2 robot that uses gaze as the feedback mechanism. We find that feedback can be useful for improving both the quality of teaching and teachers’ understanding of the robot’s capability.

Keywords: robot feedback, learning from demonstrations, algorithmic teaching

1 Introduction

If robots are to be useful to humans, they need to do more than optimize reward functions—they need to be able to figure out what reward functions should be optimized in the first place. Inverse Reinforcement Learning (IRL) [1] enables robots to infer preferences from human demonstrations. For instance, by collecting data of human drivers, we can infer a human-like driving style [2, 3]. Traditionally, IRL is applied in settings where data is collected offline from people who have no idea that a robot is supposed to learn from this data. When talking about learning preferences for the purpose of assisting people—which is arguably the goal for most robots and AI agents—there is an opportunity to explicitly involve people in teaching the robot about what they want. It can be much more effective for me to actively teach my robot how to organize my kitchen, for instance, instead of having the robot collect data of me putting things away over and over again until it eventually figures out my preferences. When I teach, I get the chance to select examples that might be especially informative to the robot—ones that effectively illustrate the core of my approach.

Unfortunately, effective teaching is tricky even when we teach other people. We have to figure out what the person knows and does not know, what teaching strategy works best for them as an individual, etc. Teaching robots is monumentally harder. We have much poorer mental models of how robots learn compared to our mental models of how humans do. What is more, when we teach humans we receive a great deal of feedback from them. One traditional way we get feedback is through tests, by asking questions to probe the learner’s understanding. More interestingly though, human students continually provide (nonverbal) feedback to the teacher *during the teaching process itself*. They look confused or bored, nod along, fidget, or gaze at various things [4, 5].

Our goal is for robots to provide similar feedback, at the same time as the teacher is providing examples. We take a step towards that in this paper, based on a simple idea: to use the states that the teacher is in as opportunities to inform the teacher about what the robot expects them to do, according to the robot’s current understanding of the task. We resort to *nonverbal cues* (e.g., gaze) as an intuitive and minimally-disruptive way for the robot to signal how it expects the teacher to

^{*}Equal contribution.

[†]Now at DeepMind.

act next. For example, imagine teaching the robot how to organize objects in a decluttering task (Fig. 1). Every time you pick up an object, the robot gazes at where it thinks the object should go. This helps you gain a better understanding of the robot’s current hypothesis about your preferences, which helps you figure out a better next example to provide. It also helps you realize when the robot has finished learning, and you can stop teaching.

Our main contributions are as follows: 1) we propose a form of feedback that robots can provide while teaching is ongoing, that consists of a prediction of the teacher’s next action along with the confidence; 2) we provide a theoretical motivation for why this form of feedback should improve teaching effectiveness, by introducing an algorithmic teaching model that takes feedback into account; and 3) we test out this feedback with real people, both in an online user study with virtual learners and in an in-person study with a real robot, where gaze is the feedback mechanism. We find that feedback can improve the teacher’s estimate of the learner’s understanding, quantifiably change the teacher’s strategy, and result in higher learner accuracy. Results with real users support our theoretical analysis. With gaze, we find that these effects are stronger when participants know explicitly about the feedback they should expect—otherwise, some participants interpret gaze differently, e.g., as acknowledgement or as purely functional.




Figure 1: The robot uses its learned preferences to anticipate the teacher’s next action (i.e., placing the object in the yellow bin), and then uses gaze to communicate its belief to the teacher.

2 Related Work

Improving quality of human input. Teaching robots via demonstration [6] is challenging; humans may have trouble providing useful demonstrations [7, 8] and knowing when to stop teaching [9]. Robots can take a more active role in learning, by asking for additional demonstrations where they are uncertain [10, 11, 12, 13] or know information is missing [14], or by asking clarification questions [15]. Robots can ask for different kinds of teacher input as well—e.g., labels, feature queries, or demonstrations—to maximize usefulness [16]. However, when robots are active learners, humans lose control over the teaching process, which can make them feel frustrated and disengaged [17, 18].

We take an orthogonal approach, in which humans maintain control of teaching, while the robot tries to be *transparent* about what it has learned. Transparency enables *algorithmic teaching*, in which the teacher’s understanding of how the learner learns enables her to select teaching examples that optimally teach this learner.³ Robot learners can be more transparent by demonstrating their current learned policy [20], allowing teachers to ask questions that probe their understanding [21, 22], or showing where they succeed and fail [9]. However, these approaches require people to stop teaching and separately test the robot. In contrast, we focus on how robots can provide feedback *at the same time* as humans are teaching—similar to the nonverbal feedback cues that human students give. This form of feedback requires minimal context switching from the teacher and does not add to the total interaction time. On the real robot, we implement this feedback as gaze.

Robot gaze. Robots can use gaze to be transparent in social interactions with humans (see Admoni and Scassellati [23] for a detailed survey). Robot gaze can also be useful in human-robot collaboration tasks, by disambiguating referenced objects [24, 25], communicating which action the robot is about to take [26], and influencing the human’s actions [27].

In this work, we also use robot gaze to communicate to humans, but specifically while the robot is *learning*. Thomaz and Breazeal [28] explored this for reinforcement learning (RL) agents: when agents use gaze to communicate what action they are about to take, humans are able to provide more informative rewards, thus speeding up learning. However, in RL the robot acts while the person observes. When robots instead learn from human demonstrations, the opposite is true, which makes the robot’s learning opaque. Taking inspiration from the RL setting, we aim at having the robot always convey what it thinks is the optimal action to take, so that the person can adjust their demonstrations appropriately.

³Since the teacher knows the task, algorithmic teaching is typically more efficient than active learning [19].

3 Nonverbal Feedback

3.1 Assumptions on robot learning algorithm

We focus on robot learners that infer *reward functions* from demonstrations via IRL [1]. The benefit of IRL is that this underlying reward function typically generalizes better across tasks, compared to directly learning a mapping from states to actions, as in behavior cloning [29, 30]. In addition, humans are naturally inclined to infer the objectives of other agents [31, 32], and thus might expect robot learners to do the same. We parameterize the reward function \mathcal{R}_θ as a linear combination of reward features $\phi(\cdot)$ with weights θ , $\mathcal{R}_\theta(s, a) = \theta^\top \phi(s, a)$, where s is the state and a is the action. There is no limitation on what these reward features can be, so this assumption is not restrictive [2].

The robot maintains a belief $b(\theta)$ over reward function parameters. We model humans as providing demonstrations (s, a) that are approximately optimal, according to a Boltzmann distribution [33]. This induces the following observation model that links human actions to the reward parameters,

$$p(a|s, \theta) \propto e^{\beta Q_\theta^*(s, a)}, \quad (1)$$

where β specifies the level of suboptimality and $Q_\theta^*(s, a)$ is the action-value function: the discounted sum of future rewards, after taking action a in state s and acting optimally thereafter with respect to \mathcal{R}_θ . As in Bayesian IRL [34], the robot uses this observation model to update its belief over θ :

$$b'(\theta) \propto p(a|s, \theta) b(\theta). \quad (2)$$

3.2 Generating feedback

We propose and investigate a form of feedback where at every step, the robot communicates what it believes the optimal action is. Intuitively, this should help human teachers understand what the robot knows and does not know as they proceed through teaching the task, enabling them to adapt what they teach and recognize when they can stop teaching. In addition, this form of feedback is minimally disruptive and allows teachers to maintain control of teaching, in contrast to tests of comprehension and active learning, respectively. We leave open the question of how this feedback should be communicated; we experiment with movement of a virtual avatar and gaze on a real robot.

Feedback target. The robot first uses its current belief over reward parameters to predict the human’s next action, using the observation model from (1):

$$\hat{a} = \arg \max_a \int_\theta p(a|s, \theta) b(\theta) d\theta. \quad (3)$$

Based on this prediction, it determines the most likely next state, i.e., where the human will go next:

$$\hat{s} = \arg \max_{s'} p(s'|s, \hat{a}). \quad (4)$$

In a deterministic environment, this is $\hat{s} = f(s, \hat{a})$, with f being the dynamics. The robot then communicates this prediction to the human teacher, for instance by gazing in the direction of \hat{s} .

Feedback speed. If the robot can control the *speed* of feedback (e.g., the speed at which it moves its end effector or gazes), it can use this to convey how *confident* it is in its prediction; slower speeds indicate lower confidence [35, 36]. Given a maximum speed of v_{max} , we thus set feedback speed as

$$v = v_{max} p(\hat{s}|s, \hat{a}) \int_\theta p(\hat{a}|s, \theta) b(\theta) d\theta. \quad (5)$$

Grounding in our experimental domain. To make this more concrete, consider the domain of decluttering: objects need to be sorted appropriately into bins, and only the human knows the correct sorting mechanism. States are locations of objects and bins, and actions place objects into bins. Every action a corresponds to a particular bin B_a , and we assume the environment is deterministic; taking action a in state s means putting the object from location s into B_a . $\phi(s, a)$ is a feature vector that consists of descriptors of object-bin match, e.g., distance, color match, shape match, etc.

Each human action teaches the robot about the relationship between the features and the reward, i.e., about the correct weights θ^* . Once the human has selected an object, the robot predicts, according to its current belief, which action \hat{a} the human will take next, and indicates $B_{\hat{a}}$, for instance via gaze. Each θ assigns different probabilities to different bins, and the robot combines all this information to compute a confidence in its estimate of bin $B_{\hat{a}}$, which it uses to adjust its feedback speed.

4 Analysis with Ideal Users: Why Do We Expect Feedback to Help?

4.1 Model of human teachers that incorporates feedback

To provide a theoretical justification for feedback, we construct potential models of human teaching that incorporate feedback, and analyze how feedback improves teachers’ decisions. These models are based on algorithmic teaching [19, 37, 38], and thus have two components: tracking the learner’s state (e.g., what does the learner know), and selecting informative demonstrations based on this.

Tracking learner state. We assume our teacher knows how the learner performs an update after every example via (2). However, the teacher is still missing information about the learner, in particular the learner’s *prior* belief b_0 , and the feature space Θ that the learner assumes—which may differ from the teacher’s. A sophisticated teacher is aware of this *uncertainty*. At every step, she has a belief over what the robot’s belief might be, that takes into account the robot’s learning updates, its feedback, and the uncertainty over which prior and feature space the robot is using.

After the first time step, this teacher computes the probability of the robot’s new belief given the robot’s feedback target x_0 and speed v_0 for the object she picked up, s_0 , and the action she took, a_0 :

$$p(b_1|s_0, a_0, x_0, v_0) = \int_{b_0, \Theta} \underbrace{p(\Theta)p(b_0|\Theta)}_{\text{teacher priors}} \underbrace{p(x_0, v_0|b_0, s_0)}_{\text{feedback generation (7)(8)}} \underbrace{p(b_1|b_0, s_0, a_0, \Theta)}_{\text{robot learning (2)}} d\Theta db_0. \quad (6)$$

Since we assume the teacher knows how the learner performs belief updates, for the robot learning distribution, $p(b_1|b_0, s_0, a_0, \Theta)$, the teacher places all probability mass on the correct new belief according to (2). For the feedback generation distribution, $p(x_0, v_0|b_0, s_0)$, we use how the robot actually generates feedback. From x_0 , the teacher recovers $\hat{a}_0 = f^{-1}(s_0, x_0)$, the action that the robot predicted, based on inverse dynamics f^{-1} . From there, the teacher computes the probability of that action under the belief b_0 :

$$p(x_0|b_0, s_0) = p(\hat{a}_0|b_0, s_0) = \int_{\theta} p(\hat{a}_0|s_0, \theta) b_0(\theta) d\theta. \quad (7)$$

For v_0 , the teacher assigns probability based on a Gaussian around her predicted speed for \hat{a}_0 :

$$p(v_0|b_0, s_0, x_0) = p(v_0|b_0, s_0, \hat{a}_0) = \mathcal{N}(v_0|v_{max} p(\hat{a}_0|b_0, s_0), v_{var}). \quad (8)$$

The teacher keeps incorporating such updates based on robot feedback for each new demonstration. In practice, we approximate this update by initializing the teacher’s priors $p(\Theta)$ and $p(b_0|\Theta)$ with sampled learner prior beliefs b_0 and possible feature spaces Θ , and iteratively updating this set of samples with the evidence, analogous to running a particle filter [39] (see Sec. 4.2 for details).

We also consider a less sophisticated teacher model, who does not account for the uncertainty over the robot’s prior and feature space. Rather than maintaining a belief over beliefs, this *iterative* teacher starts with a uniform belief over what the robot’s reward estimate might be, and updates it at every step. First, the teacher updates based on the feedback:

$$b'_1(\theta|s_0, x_0, v_0) \propto p(x_0, v_0|\theta, s_0) b_0(\theta), \quad (9)$$

with the feedback probabilities computed as in (7) and (8), but for a single θ . Next, the teacher shows a_0 and accounts for the fact that the robot will learn:

$$b_1(\theta|s_0, a_0, x_0, v_0) \propto p(a_0|\theta, s_0) b'_1(\theta). \quad (10)$$

Motivated by the “win stay, lose shift” strategy in cognitive psychology [40], iterative teachers only update if the feedback disagrees with their current belief; see Appendix A.1 for details.

Selecting examples based on current learner state. At every step, our teacher uses the tracked learner state to select the most informative example to give next. This is the example that leads to the largest increase in learner performance. The iterative teacher teaches:

$$(s_t^*, a_t^*) = \arg \max_{s_t, a_t} g(b_{t+1}(\theta|s_t, a_t)), \quad (11)$$

where $g(\cdot)$ computes the learner’s expected performance on the task (if its belief were b_{t+1}) and $b_{t+1}(\theta|s_t, a_t) \propto p(a_t|\theta, s_t) b_t(\theta)$, according to the learner’s learning update via (2).

The uncertainty-aware teacher does the same, but in expectation over the robot’s current belief and the feature space the robot uses:

$$(s_t^*, a_t^*) = \arg \max_{s_t, a_t} \mathbb{E}_{b_t, \Theta} [g(b_{t+1}(\theta|s_t, a_t))] \quad (12)$$




Figure 2: For mismatched priors, algorithmic teaching with feedback achieves higher learner performance (13). Standard error bars are computed from 100 trials.

Figure 3: With feedback, iterative and uncertainty-aware teachers (orange) can track the learner’s state accurately for most mismatch conditions. In contrast, teachers significantly overestimate learner performance when there is no feedback. Standard error bars are computed from 100 trials with ten random teaching sequences each. The y-axis is the teacher’s estimate of learner performance minus true performance (13); the x-axis is the number of teaching examples.

4.2 Impact of feedback

We now investigate how much feedback helps for our models of human teaching. If the teacher has a perfect model of the learner, then feedback is not necessary [38]. So, we focus on the kinds of mismatches that might occur—the learner might start off with a bad prior, or not know the correct feature space. For these simulated experiments, we randomly sampled N bins and O objects; each bin and object is represented by a d -dimensional feature vector. We set $N = 3$, $O = 50$, and $d = 3$.

Manipulated variables. We manipulate the teacher type: *iterative*, *uncertainty-aware*, and *random*. The *random* teacher is a baseline, that chooses random teaching examples rather than maximally informative ones. The learners vary along two factors: whether they provide feedback or not, and whether the teacher-learner mismatch is the prior (*prior mismatch*), the learner missing a feature (*feature mismatch*), or the learner learning a separate θ for each bin (*reward generalization mismatch*). For details, please refer to Appendix A.2. Note that in the latter two mismatch conditions, the learner cannot learn the task, because it is reasoning over the incorrect space of θ s.

Learner performance. We measure learning performance with expected soft classification error:

$$g(b_t(\theta)) = \mathbb{E}_{\theta \sim b_t(\theta)} \left[\frac{1}{|S|} \sum_{s \in S} p(a^*(s)|s, \theta) \right], \quad (13)$$

where $a^*(s) = \operatorname{argmax}_a \mathcal{R}_{\theta^*}(s, a)$ denotes the correct bin to place the object in.

Analysis. As one might expect, we found that the effect of feedback depends on the teacher-learner mismatch and teacher type. We found two main positive effects: feedback allows the teacher to 1) select more effective teaching examples, and 2) track the capabilities of the learner more accurately.

Feedback leads to more effective teaching for only the *prior mismatch* condition (Fig. 2). Across all mismatch conditions, uncertainty-aware teachers very quickly narrow down which learner model is correct, and can nearly-perfectly estimate the learner’s performance (Fig. 3).⁴ Although the learner cannot learn the task in the *features mismatch* and *reward generalization mismatch* conditions, there is still a benefit to the teacher estimating the learner’s performance correctly: they have a better idea of when to stop teaching (i.e., when their estimated learner performance stops increasing), and they have a more accurate estimate of the learner’s performance at the end of training. Uncertainty-aware teachers, because they maintain a belief over what the learner’s feature space might be, can also narrow down which feature(s) the learner is missing, which gives them a better idea of the learner’s capabilities for future tasks.

Summary. Our analysis on algorithmic-teaching-based human models suggests that feedback helps teachers. When the learner can learn the task, feedback makes it easier for teachers to select effective teaching examples. Feedback also improves the teacher’s estimate of the learner’s capabilities, so when the learner cannot learn the task, feedback enables users to realize this.

⁴This makes the potentially strong assumption that the uncertainty-aware teacher’s possible learner models contain the true one. This is reasonable though, in the case of *feature mismatch* for interpretable reward features, because then the teacher could just be reasoning over the power set of features.

5 Analysis in Practice: Feedback Helps

We next investigate whether real human teachers also benefit from feedback. In this section, we test the benefits of the feedback itself with virtual learners that explicitly predict the teacher’s action with varying confidence. In the next section, we test whether these benefits still exist when the feedback is realized through *gaze* on a real robot.

5.1 Design

We recruited 87 participants (ages 22-71, 41% female) on Amazon Mechanical Turk (AMT) to teach a decluttering task to an agent. Participants demonstrated correct object-bin pairings to the agent; we instructed them to give as few teaching examples as possible while still ensuring that the agent learns the correct sorting rules. The decluttering setup (Fig. 4) consisted of three numbered bins surrounded by two rings of objects, with the following sorting rules: 1) an object lying on an inner ring belongs to the closest bin, and 2) an object lying on an outer ring belongs to the bin with the same shape as it.

Users clicked on the object they wanted to demonstrate, and it would be moved into the correct bin. When feedback was activated, the learner avatar would move toward its best guess for the corresponding bin and declare that it thought the object belonged there (Fig. 4, bottom). Then the object would be moved to its bin, and the learner would acknowledge whether they were mistaken.

Manipulated variables. We manipulated the *feedback*: no feedback, full feedback that indicates the learner’s best-guess bin with variable speed corresponding to its confidence level, or partial feedback that indicates the best-guess bin with a fixed speed. We also manipulated *learner prior belief*: a uniform prior the teacher’s feature space Θ , a biased prior over Θ that heavily prefers sorting objects into their closest bin, or a uniform prior over weights in a mismatched feature space different from the teacher’s (please refer to Appendix B.2 for details). We did not conduct the 3 by 3 factorial—instead, we analyze the impact of prior condition separately from the impact of partial feedback, since we did not hypothesize any interactions there. We thus did a 2 by 3 study with full feedback and no feedback, and only tested the partial feedback on the prior belief condition.

Dependent measures. Objective— We measure the *learner performance* (13) throughout teaching. We record the *teaching sequence length* and measure proxy metrics for teaching strategy, e.g., the *proportion of objects demonstrated from each ring*. **Subjective**— We also ask open-ended and Likert scale questions about the confidence of the teacher in the learner’s understanding, their ability to track the learner’s progress, the helpfulness of feedback, and the effects of feedback on teaching. **Combined**— In addition, we compute a *mental model discrepancy* metric [18]: the human’s Likert estimate of the robot’s understanding of a sorting rule (scaled to the 0-1 range), minus the learner performance for objects classified by that rule.

Hypotheses. We hypothesize that for all learner prior conditions, **H1**: Feedback will allow the teacher to better track the learner’s progress and understanding, **H2**: The teacher will adapt their strategy according to their improved estimate of learner capabilities, and **H3**: This adjustment, enabled by feedback, will ultimately result in increased learning performance.

Subject allocation. Participants were randomly allocated between-subjects across the learner prior conditions. Full versus no feedback was within-subjects, because people may vary significantly in terms of their teaching strategies and capabilities, and we wanted to be able to make direct comparisons. We fixed the ordering of conditions: no feedback, followed by feedback. This is because the other ordering would have been too biasing—seeing the feedback first enables participants to figure out how the robot learns, and they can use that information in the no feedback condition if that happens after. Whereas with our ordering, experiencing the no feedback condition first does not inform participants’ teaching strategy, because they do not get any signal on what the robot has learned. To mitigate the risk of learning effects, we had participants first complete a practice teaching task and pass a quiz on the sorting rules.




Figure 4: AMT interface. Correct object-bin pairings are shown (top, bin 3), and learners’ feedback explicitly signals their prediction (bottom).

5.2 Results

H1: Feedback improves tracking of robot understanding. In learner prior conditions with high mental model discrepancy at the end of teaching without feedback, adding feedback reduces this discrepancy (Fig. 5). A 2 by 3 factorial repeated-measures ANOVA with feedback and prior belief as factors showed a significant interaction effect on both outer ($F(2, 84) = 19.78, p < .0001$) and inner ($F(2, 84) = 4.51, p = 0.013$) ring mental model discrepancy. The post-hoc Tukey HSD found that feedback significantly decreases discrepancy in the uniform prior condition for both the outer ($p = 0.001$) and inner ($p = 0.002$) ring rules, and in the missing feature condition for outer ($p < 0.0001$). Subjective Likert responses also support this hypothesis; see Appendix B.3 for details.

H2: Feedback changes teaching strategy. In the biased prior and the missing feature conditions, feedback leads teachers to demonstrate more outer-ring objects (Fig. 7); an ANOVA found a significant main effect for feedback on this measure ($F(1, 85) = 31.49, p < .0001$). This is reasonable, since in both conditions, the learner has trouble learning the outer ring rule. An ANOVA on the teaching sequence length found an interaction effect between the feedback and learner mismatch ($F(2, 63) = 6.613, p = .0025$). A post-hoc Tukey HSD found that the number of teaching examples is significantly higher for the missing feature case when people receive feedback—they persist in teaching the learner the outer ring rule, until they finally realize it is impossible.

H3: Feedback improves learning performance. Feedback during teaching leads to improved learner performance in the biased prior case (Fig. 6), in line with our results on ideal teacher models (Sec. 4). We ran a 2 by 3 factorial repeated-measures ANOVA on final learner performance with feedback and prior belief as factors, and number of examples as a covariate. We found a significant interaction effect between feedback and prior belief ($F(2, 1925) = 29.74, p < .0001$). A post-hoc Tukey HSD found that feedback significantly improves performance for the biased prior case, significantly decreases it for missing feature, and makes no difference for uniform prior. Since the learner cannot learn the task at all in the missing features condition, it is more important that the teacher recognizes the robot’s learning limitations—which feedback does help with.

We found no significant differences between full versus partial feedback in terms how effective teachers were (i.e., final learning performance), the proportion of demonstrated objects from the outer ring, and the number of teaching examples shown. However, learners with partial feedback do tend to overestimate learner performance for the outer ring rule (Fig. 5), possibly because they assume the learner is consistently confident in its prediction.

Summary. Overall, our results with human teachers were remarkably similar to what we saw with ideal teacher models: 1) feedback improved learning performance in the biased prior condition; and 2) although feedback did not improve performance in the uniform prior or missing feature conditions, it reduced discrepancy between the teacher’s model of the learner and the actual learner.

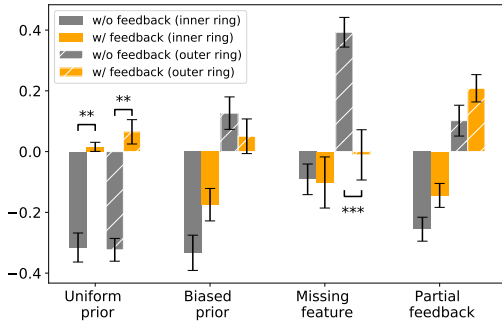


Figure 5: Mental model discrepancy (defined in Sec. 5.1) for the inner and outer ring sorting rules. *Lower magnitude is better; positive values mean that the teacher overestimates the learner capability, and negative values mean they underestimate.* (**) indicates $p < 0.005$ and (***) $p < 0.0005$.




Figure 6: Learner performance (13) versus teaching iteration in the online study (left three) and in the in-person study with a biased prior (right-most). The x-axis is the number of teaching examples.

6 Analysis in Practice: Robot Gaze

We next conducted an in-person study with a PR2 robot. In the real world, gaze cues are more natural and less disruptive than explicit communication. We did not explicitly inform users about the intent of the robot’s gaze because we hoped to evaluate whether this choice of nonverbal robot feedback is naturally interpretable—in other words, do humans naturally pick up on gaze as a feedback mechanism for action prediction?

Design. We replicated the virtual interface in a real-world decluttering task, using the same layout. Instead of a virtual learner that explicitly moves to its best-guess bin, in this study the PR2 relies on only gaze for feedback, changing its head orientation at varying speeds (Appendix B, Fig. 9).

We chose to only test the biased prior condition, because this is where feedback has the largest positive impact on learner performance—for both ideal teacher models and human teachers of virtual learners that provide feedback explicitly. We maintain the same hypotheses from the AMT study.

We recruited 17 users (ages 18-26, 47% female) from the general student population of our university—a limitation we discuss in Sec. 7. Each taught the PR2 the same task twice: first with no feedback, and then with gaze feedback. Participants were told that the robot was reset between the two tasks, so they would not assume the agent retained knowledge from previous demonstrations.

Results. Nine participants realized the gaze was related to the robot’s belief of the sorting rule. The others misunderstood the gaze as the robot either observing where the demonstrated object was selected from or acknowledging its understanding (since gazing at the bin in front resembles a nod).

We thus report findings separately for those who understood the true intent of the gaze (G1) and those who did not (G2). We ran a repeated-measures ANOVA for each dependent measure, with feedback as a factor. For G1, feedback decreases mental model discrepancy for only the inner ring rule ($F(1, 8) = 16.681, p = 0.0035$), partially supporting **H1** (Fig. 10 in Appendix). Subjective Likert responses from G1 also support **H1**; see Appendix B.4 for details. Feedback also increases the proportion of examples from the outer ring ($F(1, 8) = 9.25, p < .02$) (Fig. 7), as in the AMT study, and improves learner performance ($F(1, 8) = 9.53, p = .01$) (Fig. 6); this supports hypotheses **H2** and **H3**. None of this is true for G2. For both G1 and G2, we found no significant effect of feedback on the teaching sequence length. In addition, virtually no participants were aware of the variations in speed of the robot’s head motion, which suggests that more natural and noticeable gaze patterns (e.g., modulating acceleration) should be explored.

7 Discussion

Summary. We find that our proposed form of nonverbal robot feedback, predicting the teacher’s next action, helps improve the teacher’s effectiveness and mental model of the learner. Findings in practice echo those from our algorithmic teaching model, although we find that communicating confidence does not help significantly in practice. However, it is encouraging that despite our simple implementation of gaze, half of participants naturally interpret it as a feedback mechanism for action prediction.

Limitations and future work. Our work is limited in several ways. The main limitation is that our experiments are on a relatively simple, non-sequential task. In the real world, for more complex tasks with larger action spaces, gaze may have less communicative power; investigating effective combinations of gaze and other feedback channels is a promising future direction. In addition, due to the small and biased sample, the results of our study with gaze on the PR2 should be interpreted as trends. Further work is necessary to explore the viability of gaze as a feedback mechanism, e.g., explicitly explaining to users the purpose of gaze, or making the gaze itself more human-like.

Finally, the form of feedback we study only improves learning when the learner can learn the task, but the teacher’s model of the learner is not perfect (e.g., a mismatched prior). However, even when the task is not learnable, this form of feedback is still important in helping teachers recognize the robot’s limitations and correctly estimate its (lack of) learning.




Figure 7: Proportion of demonstrated objects from the outer ring on AMT (left) and in-person (right). (*) indicates $p < 0.05$ and (***) $p < 0.0005$.

Acknowledgments

This research was funded in part by the National Science Foundation National Robotics Initiative 2.0 (Grant #1734633) and Cyber-Physical Systems (Grant #1545126), the Center for Human-Compatible Artificial Intelligence (CHAI), and the Air Force Office of Scientific Research (Grant FA9550-17-1-0308).

References

- [1] A. Y. Ng and S. J. Russell. [Algorithms for inverse reinforcement learning](#). In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML)*, 2000.
- [2] P. Abbeel and A. Y. Ng. [Apprenticeship learning via inverse reinforcement learning](#). In *Proceedings of the Twenty-First International Conference on Machine Learning*, 2004.
- [3] S. Levine and V. Koltun. [Continuous inverse optimal control with locally optimal examples](#). In *Proceedings of the Twenty-Ninth International Conference on Machine Learning (ICML)*, 2012.
- [4] T. A. Angelo and K. P. Cross. *Classroom assessment techniques: A handbook for college teachers*. Jossey-Bass Publishers, 2nd edition, 1993.
- [5] J. M. Webb, E. M. Diana, P. Luft, E. W. Brooks, and E. L. Brennan. [Influence of pedagogical expertise and feedback on assessing student comprehension from nonverbal behavior](#). *The Journal of Educational Research*, 91(2):89–97, 1997.
- [6] B. D. Argall, S. Chernova, M. Veloso, and B. Browning. [A survey of robot learning from demonstration](#). *Robotics and Autonomous Systems*, 57(5):469–483, 2009.
- [7] M. Cakmak and A. L. Thomaz. [Optimality of human teachers for robot learners](#). In *Proceedings of the IEEE Ninth International Conference on Development and Learning (ICDL)*, 2010.
- [8] F. Khan, B. Mutlu, and J. Zhu. [How do humans teach: On curriculum learning and teaching dimension](#). In *Advances in Neural Information Processing Systems (NeurIPS) 24*, 2011.
- [9] A. Sena, Y. Zhao, and M. J. Howard. [Teaching human teachers to teach robot learners](#). In *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [10] T. Inamura, M. Inaba, and H. Inoue. [Acquisition of probabilistic behavior decision model based on the interactive teaching method](#). In *Proceedings of the Ninth International Conference on Advanced Robotics (ICAR)*, 1999.
- [11] A. P. Shon, D. Verma, and R. P. N. Rao. [Active imitation learning](#). In *Proceedings of the 22nd National Conference on Artificial Intelligence (AAAI)*, 2007.
- [12] S. Chernova and M. Veloso. [Interactive policy learning through confidence-based autonomy](#). *Journal of Artificial Intelligence Research*, 34(1):1–25, 2009.
- [13] D. Silver, J. A. Bagnell, and A. Stentz. [Active learning from demonstration for robust autonomous navigation](#). In *IEEE International Conference on Robotics and Automation (ICRA)*, 2012.
- [14] J. Miura, K. Iwase, and Y. Shirai. [Interactive teaching of a mobile robot](#). In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA)*, 2005.
- [15] D. Whitney, E. Rosen, J. MacGlashan, L. L. S. Wong, and S. Tellex. [Reducing errors in object-fetching interactions through social feedback](#). In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [16] M. Cakmak and A. L. Thomaz. [Designing robot learners that ask good questions](#). In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2012.
- [17] C. Chao, M. Cakmak, and A. L. Thomaz. [Transparent active learning for robots](#). In *Proceedings of the Fifth ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2010.
- [18] M. Cakmak, C. Chao, and A. L. Thomaz. [Designing interactions for robot active learners](#). *IEEE Transactions on Autonomous Mental Development*, 2(2):108–118, June 2010.
- [19] X. Zhu. [Machine teaching: An inverse problem to machine learning and an approach toward optimal education](#). In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

- [20] S. Calinon and A. Billard. [Active teaching in robot programming by demonstration](#). In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2007.
- [21] L. Steels and F. Kaplan. [AIBO's first words: The social learning of language and meaning](#). *Evolution of Communication*, 4(1):3–32, 2001.
- [22] G. H. Lim, M. Oliveira, V. Mokhtari, S. H. Kasaei, A. Chauhan, L. S. Lopes, and A. M. Tomé. [Interactive teaching and experience extraction for learning about objects and robot activities](#). In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2014.
- [23] H. Admoni and B. Scassellati. [Social eye gaze in human-robot interaction: A review](#). *Journal of Human-Robot Interaction*, 6(1):25–63, 2017.
- [24] H. Admoni, T. Weng, and B. Scassellati. [Modeling communicative behaviors for object references in human-robot interaction](#). In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2016.
- [25] J.-D. Boucher, U. Pattacini, A. Lelong, G. Bailly, F. Elisei, S. Fagel, P. F. Dominey, and J. Ventre-Dominey. [I reach faster when I see you look: Gaze effects in human-human and human-robot face-to-face cooperation](#). *Frontiers in Neurorobotics*, 6(3), 2012.
- [26] A. Moon, D. M. Troniak, B. Gleeson, M. K. X. J. Pan, M. Zheng, B. A. Blumer, K. MacLean, and E. A. Croft. [Meet me where I'm gazing: How shared attention gaze affects human-robot handover timing](#). In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2014.
- [27] H. Admoni, A. Dragan, S. Srinivasa, and B. Scassellati. [Deliberate delays during robot-to-human handovers improve compliance with gaze communication](#). In *International Conference on Human-Robot Interaction (HRI)*, 2014.
- [28] A. L. Thomaz and C. Breazeal. [Teachable robots: Understanding human teaching behavior to build more effective robot learners](#). *Artificial Intelligence*, 172(6-7):716–737, 2008.
- [29] D. A. Pomerleau. [Efficient training of artificial neural networks for autonomous navigation](#). *Neural Computation*, 3(1):88–97, March 1991.
- [30] S. Ross, G. J. Gordon, and J. A. Bagnell. [A reduction of imitation learning and structured prediction to no-regret online learning](#). In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- [31] C. L. Baker, R. Saxe, and J. B. Tenenbaum. [Action understanding as inverse planning](#). *Cognition*, 113(3):329–349, 2009.
- [32] J. Jara-Ettinger, H. Gwen, L. E. Schulz, and J. B. Tenenbaum. [The naïve utility calculus: Computational principles underlying commonsense psychology](#). *Trends in Cognitive Sciences*, 20(8):589–604, 2016.
- [33] P. A. Ortega and A. A. Stocker. [Human decision-making under limited time](#). In *Advances in Neural Information Processing Systems (NeurIPS) 29*, 2016.
- [34] D. Ramachandran and E. Amir. [Bayesian inverse reinforcement learning](#). In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.
- [35] J. Hough and D. Schlagen. [It's not what you do, it's how you do it: Grounding uncertainty for a simple robot](#). In *Proceedings of the Twelfth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2017.
- [36] A. Zhou, D. Hadfield-Menell, A. Nagabandi, and A. D. Dragan. [Expressive robot motion timing](#). In *Proceedings of the Twelfth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2017.
- [37] S. A. Goldman and M. J. Kearns. [On the complexity of teaching](#). *Journal of Computer and System Sciences*, 50(1):20–31, 1995.
- [38] F. J. Balbach and T. Zeugmann. [Recent developments in algorithmic teaching](#). In *Proceedings of the 3rd International Conference on Language and Automata Theory and Applications (LATA)*, 2009.
- [39] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. [Novel approach to nonlinear/non-Gaussian Bayesian state estimation](#). *IEE Proceedings F - Radar and Signal Processing*, 140(2):107–113, April 1993.
- [40] M. Nowak and K. Sigmund. [A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game](#). *Nature*, 364:56–58, 1993.

A Analysis with Ideal Users

A.1 Model of iterative teachers

Motivated by the “win stay, lose shift” strategy in cognitive psychology [40], we model iterative teachers as only updating their belief if the feedback disagrees with their current belief in the learner’s state. In other words, if the learner predicts the bin that the teacher’s maximum a posteriori (MAP) estimate of the learner’s θ deems most likely, and the corresponding velocity for this θ_{MAP} is within $0.05 * v_{\text{max}}$ of the learner’s feedback speed, then the teacher does not update based on this feedback. Thus, if the teacher has a perfect model of how the learner learns, then whether the learner provides feedback or not, the teacher would maintain the same belief over the learner’s hypothesis, and thus teach in the exact same way. If we did not make integrating feedback conditional, then this teacher would potentially teach *worse* in this situation, since the additional feedback-based updates cause the teacher’s estimate to diverge from the true one.

A.2 Implementation details

Learner models. In the *prior mismatch* condition, we bias the learner toward a randomly-selected θ' by setting the prior as the softmax of the distance between θ s:

$$p(\theta) \propto e^{\beta' \theta^\top \theta'}, \quad (14)$$

with $\beta' = 50$ for a fairly strong bias. In the *feature mismatch* condition, the learner ignores the last feature dimension. In other words, the set of θ s that the learner considers all have $\theta_d = 0$. In the *reward generalization mismatch* condition, the learner learns a separate θ for each bin, i.e., reward features are computed in terms of only the object s , as in $\phi(s)$. So the space of θ s the learner is considering now has dimensions $B * d$.

For learning updates, we set $\beta = 20$ for a learner that assumes demonstrations are close to optimal. For feedback, we set $v_{\text{max}} = 1.0$ and enforce a minimum speed of 0.05. We sampled 1000 θ s to approximate the learner’s belief.

Teacher models. In the *prior mismatch* condition, the uncertainty-aware teacher considers $P + 1$ possible learner models, one of which is the default (i.e., uniform prior b_0); the other P are biased towards different θ s, including the one that the learner is actually biased towards. In the *feature mismatch* condition, the uncertainty-aware teacher considers $d + 1$ possible learner models: one with all features, and d that are each missing a different one of the d features. Finally, in the *reward generalization mismatch* condition, the uncertainty-aware teacher considers two possible learner models, one of which has reward features $\phi(s, a)$ that depend on both the object and bin, and the other of which has reward features $\phi(s)$ that depend only on the object. We assume uncertainty-aware teachers start off believing that the learner likely has a uniform prior b_0 , considers all d features, and uses reward features $\phi(s, a)$ —so we set the teachers’ priors (i.e., $p(b_0|\Theta)$ and $p(\Theta)$) to place equal probability on all of the models, except ten times more probability on this most-likely learner model. For teacher updates, we assume $v_{\text{var}} = 0.0025$.

B Analysis in Practice

B.1 Participant consent and compensation

Both the AMT and in-person user studies, along with their corresponding consent forms, received IRB approval. Participants were paid \$3.75 for the AMT study, which lasted around 15 minutes, and participants were given a \$10 Amazon gift card for the in-person study, which lasted just under 20 minutes.

B.2 Implementation details

We generated 1024 θ s for the learner to reason over, including those that corresponded to conceptually intuitive sorting rules (e.g., objects belong to bins with their same shape). In the uniform prior case, the belief was uniform over these θ s, and in the biased prior case, the learner heavily preferred sorting all objects into their closest bin. In the mismatched features condition, we removed the shape dimension feature in the learner’s belief space.




Figure 8: Likert scale responses from the AMT study, in which users taught learners with and without feedback. (*) indicates $p < 0.05$, (**) indicates $p < 0.005$ and (***) indicates $p < 0.0005$. Labels correspond to the Likert questions in Table 1.

B.3 Analysis of Likert responses from AMT study

For the AMT study, we ran a two-way repeated measures ANOVA on each of the subjective measures (i.e., Likert questions) with the feedback and prior belief as factors. The full details are reported in Table 1. For clarity, in the following analysis we will refer to the questions by their corresponding labels; full questions can be found in the first column of Table 1.

Tracking robot understanding. In the uniform prior condition, when the learner is able to learn both the inner and outer ring sorting rules, participants’ subjective Likert responses indicate that when they were given feedback, they were significantly more confident that the learner learned correctly (Fig. 8, *inner* and *outer*). A post-hoc Tukey HSD found a significant difference between feedback and no feedback for both *inner* ($p = 0.004$) and *outer* ($p = 0.004$).

Similarly, in the missing feature condition, when the learner is not able to learn the outer ring rule, participants who received feedback were more aware of this; a post-hoc Tukey HSD found a significant difference between feedback and no feedback for *outer* ($p = 0.0003$).

Across all prior belief conditions, participants found it overwhelmingly easier to recognize what the learner *does not* understand; the ANOVA found a significant main effect for feedback for *doesn't*. Participants also generally found it easier to recognize what a learner *knows*; a post-hoc Tukey HSD found a significant difference between feedback and no feedback for the uniform prior ($p < 0.0001$) and partial feedback ($p < 0.0001$) conditions.

User satisfaction. For *happy*, a post-hoc Tukey HSD found a significant difference between feedback and no feedback for the missing feature condition ($p = 0.0011$). This makes sense, because when feedback was given in the missing feature condition, participants were clearly aware that the robot was unable to learn the task, despite their efforts to teach it.

Full versus partial feedback. An ANOVA on each of the subjective measures with the feedback level as a factor (full feedback versus partial feedback) found significant differences between the partial and full feedback condition for only the *stop* ($F(1, 47) = 8.43, p = 0.0059$) and *happy* ($F(1, 47) = 5.57, p = 0.02$) questions. The average for the partial feedback condition was higher in both cases. We hypothesize that this is because participants found the partial feedback (with constant speed) to be easier to reason about. This suggests that there is a need for investigating other mechanisms for expressing feedback confidence.

Statement	Effect	F-score	p-value
Inner: “[Learner] correctly learned the rule that if an object lies in an inner ring, it belongs to the closest bin”	interaction	$F(3, 83) = 7.99$	$p < 0.0001$
Outer: “[Learner] correctly learned the rule that if an object lies in an outer ring, it belongs to the bin with the same shape”	interaction	$F(3, 83) = 14.12$	$p < 0.0001$
Informative: “I found it easy to choose informative teaching examples for [Learner]”	none	—	—
Stop: “It was easy to know when to stop teaching [Learner]”	none	—	—
Knows: “It was easy to tell what [Learner] knows about the task rules”	interaction	$F(3, 83) = 3.89$	$p = 0.01$
Doesn’t: “It was easy to tell what [Learner] <i>doesn’t</i> know about the task rules”	interaction	$F(1, 83) = 87.7$	$p < 0.0001$
Happy: “I would be happy to teach [Learner] again”	interaction	$F(3, 83) = 6.14$	$p = 0.0008$

Table 1: ANOVA Results for AMT Study Likert Questions.



Figure 9: Real-world experiment flow: the participant a) selects the object to demonstrate, b) picks it up and shows it to the PR2, c) observes the PR2’s gaze feedback, and d) places the object in the correct bin.

B.4 Analysis of Likert responses from in-person study

For the in-person study with the PR2 robot, we ran a repeated-measures ANOVA on each of the subjective measures (i.e., Likert questions) with feedback as a factor. These questions are the same as for the AMT study. The full details are reported in Table 2.

Tracking robot understanding. Recall that in this study, the PR2 robot started with a biased prior over weights in Θ , that heavily favors sorting all objects into their closest bin. This biased prior agrees with the inner rule, but contradicts the outer rule. When participants received gaze feedback from the robot, those who understood the true intent of the robot’s gaze (G1) better recognized that it correctly learned the *inner* rule ($F(1, 8) = 16.66, p = 0.0035$), compared to when there was no feedback. Participants in G1 also thought it was easier to give *informative* teaching examples ($F(1, 8) = 12.57, p = 0.0076$), recognize when to *stop* teaching ($F(1, 8) = 14.59, p = 0.0051$), and recognize what the robot *knows* ($F(1, 8) = 139.13, p < 0.0001$) and *doesn’t* know ($F(1, 8) = 144.00, p < 0.0001$) when the robot provided gaze feedback.

Even participants who misunderstood the intent of the robot’s gaze (G2) thought it was easier to recognize when to *stop* teaching ($F(1, 7) = 6.66, p = 0.036$) and recognize what the robot *knows* ($F(1, 7) = 5.64, p = 0.049$) when the robot provided gaze feedback.

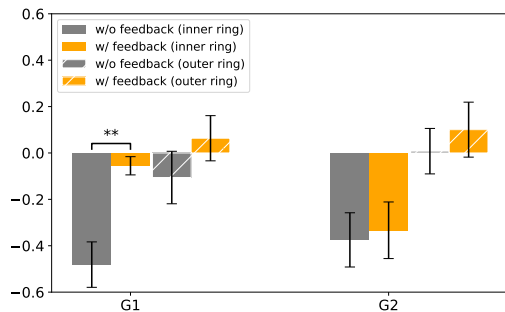


Figure 10: Mental model discrepancy (defined in Sec. 5.1) for the inner and outer ring sorting rules. **Lower magnitude is better; positive values mean that the teacher overestimates the learner capability, and negative values mean they underestimate.** (**) indicates $p < 0.005$.




Figure 11: Likert scale responses from the in-person study, in which users taught learners with and without feedback. (*) indicates $p < 0.05$, (**) indicates $p < 0.005$ and (***) indicates $p < 0.0005$. Labels correspond to the Likert questions in Table 2.

Statement	Group	F-score	p-value
Inner: “[Learner] correctly learned the rule that if an object lies in an inner ring, it belongs to the closest bin”	G1	$F(1, 8) = 16.66$	$p = 0.0035$
	G2	—	—
Outer: “[Learner] correctly learned the rule that if an object lies in an outer ring, it belongs to the bin with the same shape”	G1	—	—
	G2	—	—
Informative: “I found it easy to choose informative teaching examples for [Learner]”	G1	$F(1, 8) = 12.57$	$p = 0.0076$
	G2	—	—
Stop: “It was easy to know when to stop teaching [Learner]”	G1	$F(1, 8) = 14.59$	$p = 0.0051$
	G2	$F(1, 7) = 6.66$	$p = 0.036$
Knows: “It was easy to tell what [Learner] knows about the task rules”	G1	$F(1, 8) = 139.13$	$p < 0.0001$
	G2	$F(1, 7) = 5.64$	$p = 0.049$
Doesn't: “It was easy to tell what [Learner] <i>doesn't</i> know about the task rules”	G1	$F(1, 8) = 144.00$	$p < 0.0001$
	G2	—	—
Happy: “I would be happy to teach [Learner] again”	G1	$F(1, 8) = 10.73$	$p = 0.011$
	G2	—	—

Table 2: ANOVA Results for In-Person Study Likert Questions.

User satisfaction. Participants in G1 were more *happy* to teach the robot in the future, when it provided gaze feedback ($F(1, 8) = 10.73$, $p = 0.011$).